

NONCROSSING PARTITIONS IN SURPRISING LOCATIONS

JON MCCAMMOND

1. INTRODUCTION.

Certain mathematical structures make a habit of reoccurring in the most diverse list of settings. Some obvious examples exhibiting this intrusive type of behavior include the Fibonacci numbers, the Catalan numbers, the quaternions, and the modular group. In this article, the focus is on a lesser known example: the noncrossing partition lattice. The focus of the article is a gentle introduction to the lattice itself in three of its many guises: as a way to encode parking functions, as a key part of the foundations of noncommutative probability, and as a building block for a contractible space acted on by a braid group. Since this article is aimed primarily at nonspecialists, each area is briefly introduced along the way.

The noncrossing partition lattice is a relative newcomer to the mathematical world. First defined and studied by Germain Kreweras in 1972 [33], it caught the imagination of combinatorialists beginning in the 1980s [20], [21], [22], [23], [29], [37], [39], [40], [45], and has come to be regarded as one of the standard objects in the field. In recent years it has also played a role in areas as diverse as low-dimensional topology and geometric group theory [9], [12], [13], [31], [32] as well as the noncommutative version of probability [2], [3], [35], [41], [42], [43], [49], [50]. Due no doubt to its recent vintage, it is less well-known to the mathematical community at large than perhaps it deserves to be, but hopefully this short paper will help to remedy this state of affairs.

2. A MOTIVATING EXAMPLE.

Before launching into a discussion of the noncrossing partition lattice itself, we quickly consider a motivating example: the Catalan numbers. The Catalan numbers are a favorite pastime of many amateur (and professional) mathematicians. In addition, they also have a connection with the noncrossing partition lattice (Theorem 3.1).

Example 2.1 (Catalan numbers). The Catalan numbers are the numbers C_n given by

$$C_n = \frac{1}{n+1} \binom{2n}{n},$$

and they have a number of different interpretations. See, for example, Richard Stanley's list of more than one hundred distinct ways in which this sequence arises [44]. Some of the most common interpretations are as the number of triangulations of an $(n+2)$ -sided polygon (illustrated in Figure 1), as the number of binary parenthesizations of a string of $n+1$ letters (Figure 2), or as the number of rooted trivalent plane trees with $2n+2$ vertices (Figure 3). With these examples to whet




FIGURE 1. The fourteen triangulations of a hexagon.

$$\begin{array}{ll}
 (1(2(3(45)))) & (1(2((34)5))) \\
 ((1(23)(45))) & (1((2(34))5)) \\
 (1(((23)4)5)) & ((12)(3(45))) \\
 ((12)((34)5)) & ((1(23))(45)) \\
 (((12)(34))5) & ((1((23)4))5) \\
 (((12)3)(45)) & (((12)(34))5) \\
 (((1(23))4)5) & (((((12)3)4)5))
 \end{array}$$

FIGURE 2. The fourteen ways to associate five numbers.




FIGURE 3. The fourteen rooted trivalent plane trees with six leaves and ten vertices.

the reader's appetite, we direct the interested reader to [44] and [47, Exercise 6.19], and we continue on to a description of noncrossing partitions.

3. NONCROSSING PARTITIONS.

We are now ready to define a noncrossing partition. Following traditional combinatorial practice we use $[n]$ to denote the set $\{1, \dots, n\}$.

Noncrossing partitions. Recall that a *partition* of a set is a collection of pairwise disjoint subsets whose union is the entire set and that the subsets in the collection are called *blocks*. A *noncrossing partition* σ is a partition of the vertices of a regular n -gon (labeled by the set $[n]$) so that the convex hulls of its blocks are pairwise disjoint. Figure 4 illustrates the noncrossing partition $\{\{1, 4, 5\}, \{2, 3\}, \{6, 8\}, \{7\}\}$. The partition $\{\{1, 4, 6\}, \{2, 3\}, \{5, 8\}, \{7\}\}$ would be crossing.

Given partitions σ and τ of $[n]$ we say that $\sigma < \tau$ if each block of σ is contained in a block of τ . This ordering on the set of all partitions of $[n]$ defines a partially ordered set called the *partition lattice* and is usually denoted Π_n . When restricted to the set of noncrossing partitions on $[n]$, it is called the *noncrossing partition lattice* and denoted NC_n . The poset Π_4 is shown in Figure 5. For $n = 4$, the only difference between the two posets is the partition $\{\{1, 3\}, \{2, 4\}\}$, which is not noncrossing.

The poset of noncrossing partitions has a number of nice combinatorial properties that we record in the following theorem:





FIGURE 4. A noncrossing partition of the set [8].

FIGURE 5. The figure shows the partition lattice for $n = 4$. If the vertex surrounded by a dashed line is removed, the result is the noncrossing partition lattice for $n = 4$.

Theorem 3.1. *For each n the poset NC_n is a graded, bounded lattice with Catalan many elements (C_n to be exact). In addition, it is self-dual and locally self-dual.*

The fact that the number of noncrossing partitions is a Catalan number is part of the lore of the Catalan numbers, and we refer the reader again to [44] and [47]. For the other properties we now review their definitions.

A poset is *bounded* if it has both a minimum element and a maximum element. For the noncrossing partition lattice, the discrete partition (i.e., the one in which each block contains a single element) and the partition with a single block fulfill these roles. A *chain* in a poset is a subset in which any two elements are comparable, its *length* is one less than the size of this subset, and a *maximal chain* is a chain that is not properly contained in any larger chain. A poset in which any two maximal chains have the same length is called *graded*. In any graded bounded poset there

is a height function that keeps track of the level in which elements are contained. The level of an element x can be defined as the number of elements strictly below x in any maximal chain containing x . In Π_n or NC_n , for example, the height of a partition is n minus the number of blocks. For later use, we also note that if $\sigma < \tau$ in Π_n or NC_n and if their heights differ by one, then there are blocks B and B' in σ whose union is a block in τ . Moreover, all other blocks in σ and τ are identical. In this situation we say that τ *covers* σ . A chain in which each element covers the previous one is called a *saturated chain*.

A poset is a *lattice* if each pair of elements has a least upper bound and a greatest lower bound. The greatest lower bound $\sigma \wedge \tau$ of partitions σ and τ of $[n]$ is simply the largest refinement of the two partitions. In other words, define i and j to be in the same block of $\sigma \wedge \tau$ if and only if they lie in the same blocks in both σ and τ . It is now easy to see that this is a lower bound for σ and τ and that this is greater than any other lower bound. To find the least upper bound $\sigma \vee \tau$ of σ and τ from NC_n , superimpose the convex hulls of all the blocks for σ and for τ and then take the convex hulls of the connected components that result.

Finally, a poset is *self-dual* if there is an order-reversing bijection from it to itself, and it is *locally self-dual* if this is true for each of its intervals. Recall that an interval $[x, y]$ in a poset is simply the subposet containing all the elements greater than or equal to x and less than or equal to y . Since this property is easier to establish once we make the connection to the symmetric groups, we postpone its proof until the next section.

4. SYMMETRIC GROUPS.

Before connecting the noncrossing partition lattice with low-dimensional topology and noncommutative probability, it will be helpful to establish first its close connection with the symmetric group.

Symmetric groups. A *permutation* of a set X is a bijection from X to itself. We use S_n to signify the group of all permutations of the set $[n]$ under function composition. (We refer to S_n as the *symmetric group* on n elements.) There are two natural generating sets for a group of permutations. If the underlying set is unordered, then the most natural generating set is the set of permutations that interchange two elements and leave the rest fixed (usually called *transpositions* or *two-cycles*). If, on the other hand, the underlying set itself has a natural linear ordering, as is the case for $[n]$, then the smaller generating set using only adjacent transpositions $(i, i + 1)$ is often preferred.

One of the more mysterious properties of permutations for students in an abstract algebra class is the fact that they are naturally classified as either even or odd. Since many of the standard proofs of this are unenlightening and since there is an elementary geometric proof, we digress slightly to present it. The following result is, of course, the heart of the matter:

Theorem 4.1. *Every product of transpositions that equals the identity permutation has an even number of factors.*

Proof. We first prove the theorem in the case where all of the transpositions involved are adjacent ones. Given a sequence of adjacent transpositions, we can draw a set of curves as illustrated in Figure 6. The intersections correspond to the transpositions




FIGURE 6. Illustration of the proof of Theorem 4.1.

as follows. Arrange the transpositions from top to bottom in the order they are applied. In the example, the first transposition is $(3, 4)$ and the curves on the left are drawn so that highest crossing occurs between the strands that are third and fourth from the left. The next transposition is $(2, 3)$ so the curves are drawn so that the next highest crossing occurs between the strands that are currently second and third from the left. This procedure can be used to convert any sequence of adjacent transpositions into smooth descending curves. (This procedure can also be reversed for smooth descending curves in general position.)

If the product is indeed the identity permutation, then the curve that starts at position i also ends at position i . Call this the i th curve. Now we simply change perspectives. Each intersection is an intersection between two curves, say the i th and j th curves. Since the i th and j th curves originally occur in a particular order and they also return to the same order, these two curves must intersect an even number of times. Adding up the intersections according to the curves involved gives the result. To convert this to a result about products of arbitrary transpositions, it remains only to note that every transposition can be written as a product of an odd number of adjacent transpositions, so products over the larger generating set can be converted to products over the smaller generating set without changing the parity of the factorization. \square

Using the usual trick of rewriting a pair of factorizations of a permutation as a single factorization of the identity, the following corollary is immediate:

Corollary 4.2. *The parity of a factorization of a permutation into transpositions is independent of the factorization chosen.*

Returning our attention to the noncrossing partition lattice, we find that NC_n is closely connected with the factorizations of an n -cycle into transpositions. First, we introduce some definitions.

Minimal factorizations. A factorization of a permutation into transpositions is called *minimal* if it has minimal length (i.e., the smallest number of factors) among all such factorizations. We can then define an ordering on the set of permutations by declaring that $\sigma < \tau$ if there is a minimal factorization of τ with a “prefix” that is a minimal factorization of σ . In other words, there should be a minimal factorization $t_1 \circ t_2 \circ \dots \circ t_k = \tau$ such that $t_1 \circ t_2 \circ \dots \circ t_\ell = \sigma$ with $\ell \leq k$.

One quick note about multiplication. There are two natural conventions for multiplying permutations: functional notation and algebraist notation. We use the algebraic convention throughout so that $(1, 2)(1, 3)$ is $(1, 2, 3)$ rather than $(1, 3, 2)$. Because of all the symmetries of the objects under consideration this is only of minor importance, but it means that we need to write “prefix” rather than “suffix” in the foregoing definition.

Lemma 4.3. *The poset of permutations less than or equal to the n -cycle $(1, 2, 3, \dots, n)$ is isomorphic with the noncrossing partition lattice NC_n .*

Sketch of proof. Instead of giving a complete proof we write down the isomorphism and omit the details. The interested reader can find a complete proof in [12]. Given a noncrossing partition σ , we convert it into a permutation $\pi(\sigma)$ by writing each block as a disjoint cycle. The order in which the elements occur is the natural linear order, or stated more geometrically, we read the labels of each convex hull clockwise. To illustrate, the noncrossing partition σ illustrated in Figure 4 corresponds to the permutation $\pi(\sigma) = (1, 4, 5)(2, 3)(6, 8)$. Showing that the permutations corresponding to noncrossing partitions are prefixes of reduced factorizations of the n -cycle $(1, 2, \dots, n)$ is relatively easy, as is the correspondence between the two orderings. The only slightly tricky part is showing that this map is onto. \square

What Lemma 4.3 proves, in essence, is that the graph whose vertices are elements of NC_n and whose edges are its covering relations corresponds to a portion of the right Cayley graph of S_n with respect to the set of all transpositions. We remind readers that the *Cayley graph* of a group G with respect to a generating set A is a directed graph with vertices labeled by the elements of G and edges indexed by the set $G \times A$, where the edge (g, a) connects g to $g \cdot a$ [27]. Because of this identification, not only does every element of NC_n have a permutation assigned to it, but if τ covers σ then this edge is labeled by a transposition (equal to $\pi(\sigma)^{-1}\pi(\tau)$). We have illustrated this labeling for NC_3 in Figure 7. Notice that the Cayley graph interpretation also ensures that the sequence of edge labels in a maximal chain, read from the bottom to the top, multiply together to give $(1, 2, \dots, n)$ and that these correspond exactly to its minimal factorizations. For example, the minimal factorizations of $(1, 2, 3)$ into transpositions (in algebraist notation) are $(1, 2)(1, 3) = (1, 3)(2, 3) = (2, 3)(1, 2) = (1, 2, 3)$ and these are the labels on the three maximal chains seen in Figure 7.

More generally, multiplying together the sequence of transpositions labeling the edges in a saturated chain connecting σ to τ yields the permutation $\pi(\sigma)^{-1}\pi(\tau)$. Using Lemma 4.3 it is now easy to establish the following:

Proposition 4.4. *If $\rho < \sigma < \tau$ in NC_n and $(\alpha_1, \alpha_2, \dots, \alpha_k)$ is a sequence of transpositions labeling a saturated chain from ρ to σ , then there is a unique element σ' in NC_n and a saturated chain from σ' to τ with the exact same sequence of labels. Similarly, if $(\beta_1, \beta_2, \dots, \beta_k)$ is a sequence of transpositions labeling a saturated chain from σ to τ , then there is a unique element σ'' in NC_n and a saturated chain from ρ to σ'' with this same sequence of labels.*

Proof. This statement is actually a consequence of the observation that the set of transpositions in S_n is closed under conjugation. Thus, if α and β are single transpositions and $\alpha\beta$ is a minimal factorization, then there is another minimal factorization $\gamma\alpha$, where γ is the transposition $\alpha\beta\alpha = (\alpha\beta\alpha^{-1})$. Iterating this idea




FIGURE 7. The noncrossing partition lattice for $n = 3$ with edge and vertex labels.

allows us to move labeled chains up and down as much as we want. Finally, the partitions σ' and σ'' must be unique since we know exactly the permutations to which they correspond under the identification with the Cayley graph. \square

The locally self-dual property is now almost immediate:


Proof that NC_n is (locally) self-dual. Given σ and τ in NC_n such that $\sigma < \tau$, let P be the poset of noncrossing partitions between them. Define a map $f : P \rightarrow P$ so that for each ρ in P the labels on a saturated chain from σ to ρ are also the labels on a saturated chain from $f(\rho)$ to τ . By Proposition 4.4 a unique such element exists, and it is easy to check that f is a well-defined order reversing isomorphism from P and P . \square

Finally, we note that the symmetric groups are examples of a broader class of groups called *finite reflection groups* or *finite Coxeter groups*. In the same way that the noncrossing partition lattice is closely connected with the symmetric group, there is an entire series of lattices, one for each finite Coxeter group [7], [9], [13]. Each of these general noncrossing partition lattices is graded, bounded, self-dual, and locally self-dual, and the general proofs are essentially the same as the ones given here. Because of these elegant patterns, the number of elements in these additional lattices have come to be called *generalized Catalan numbers*. In the course of the article we occasionally comment on properties that extend to these additional situations.

5. BRAID GROUPS.

We are now ready to establish our first connection: noncrossing partitions and the braid groups. The braid groups are related to many areas of mathematics including mathematical physics, quantum groups, von Neumann algebras, and, not too surprisingly, three-manifold topology (see, for example [26], [28], [30], [34], [36], or [38]). The surprise is that they are also intimately related to noncrossing partitions.

Roughly stated, a *braid on n strings* keeps track of how n strings can be twisted in space so that each strand is a smooth embedded monotonically decreasing curve in

FIGURE 8. An element of B_4 .

\mathbb{R}^3 (i.e., its partial with respect to z is always strictly negative). Various conventions need to be established, such as that strands must start and end in some standardized configuration, that strands cannot intersect, and that perturbations of a braid that maintain these conventions are considered to be the same element. The collection of all braids is turned into the braid group B_n once multiplication is defined by attaching one braid to the top of another (see Figure 8 for a typical element of B_4). Observe that there is a group homomorphism from B_n onto S_n that simply forgets which way crossings took place and merely records the permutation of the strings (compare, for example, Figures 8 and 6). The way the noncrossing partition lattice enters the picture is through the structure of its order complex.

Each partially ordered set P can be turned into a simplicial complex in a very simple fashion. The vertices of the complex are labeled by the elements of the poset, and we add a simplex corresponding to a set of vertices if and only if the elements that label them form a chain in P . The result is called the *order complex* of P (or its *geometric realization*).

Let $\Delta(NC_n)$ denote the order complex of NC_n . Notice that the one-cells in $\Delta(NC_n)$ correspond to two-element chains in NC_n . In other words, they correspond to pairs of noncrossing partitions σ and τ with $\sigma < \tau$. We can carry over the labeling from NC_n to the one-cells in $\Delta(NC_n)$ as follows: we label the oriented edge from σ to τ in $\Delta(NC_n)$ by the permutation $\pi(\sigma)^{-1}\pi(\tau)$. Our final step is identify certain simplices in $\Delta(NC_n)$ with each other. The rule is that two simplices are identified if and only if they have identically labeled, oriented one-skeletons (and they are identified so that these labels match, of course). The resulting quotient is a complex we call the *Brady-Krammer complex* BK_n , since it was discovered independently by Tom Brady and Daan Krammer. What Brady and Krammer proved in [12] and [31], respectively, was that this procedure results in a complex whose fundamental group is the braid group B_n and whose universal cover is contractible. In other words, they proved the following result:

Theorem 5.1 (Brady, Krammer). *The complex BK_n is an Eilenberg-Maclane space for the braid group B_n .*

As in the previous section, the procedure described extends naturally to the general noncrossing partition lattices associated with the other finite reflection groups. In each case the resulting complex is an Eilenberg-Maclane space for a group, and this group is related to the finite reflection group in the same way that the braid group is related to the symmetric group. These other groups are called *finite-type*

Artin groups. See [9] or [13] for a proof of this extension or [10], [15], [16], or [17] for more about finite-type Artin groups.

6. PARKING FUNCTIONS.

For our second illustration we shift to a classic problem from combinatorics. Imagine a sequence of n cars entering a one-way street one at a time with n parking spots available, as shown in Figure 9. Each driver has a preferred parking spot and attempts to park there first. Failing that, he or she parks in the next available space. If any of the drivers is forced out of the street, then this sequence of preferences has failed. We refer to any sequence of preferences that enables all n of the drivers to park successfully as a *parking function*.




FIGURE 9. A one-way street with 7 parking spots.

As an initial observation, it is easy to see that if two or more drivers prefer the final parking space, the sequence fails. Similarly, if three or more drivers prefer either of the last two spaces, the sequence fails, etc. Perhaps surprisingly, these are the only restrictions. Thus, an equivalent definition of a parking function would describe it as a sequence (a_1, \dots, a_n) of positive integers whose rearrangement (and relabeling) as a nondecreasing sequence $b_1 \leq b_2 \leq \dots \leq b_n$ satisfies the inequalities $b_i \leq i$ for each i in $[n]$. From this altered definition it is not hard to show that the number of parking functions is $(n+1)^{n-1}$. Combinatorialists, of course, recognize this as the number that counts labeled rooted trees on $[n]$ (or, equivalently, the number of acyclic functions on $[n]$). In order to see the connection with the noncrossing partition lattice, consider the following definition due to Richard Stanley.

Suppose that τ covers σ in NC_{n+1} , that B and B' are the two blocks of σ that combine to form a block in τ , and without loss of generality, that $\min B < \min B'$. We define a new label on this covering relation by the largest element of B that is below each element of B' . Using this edge-labeling Stanley was able to show the following [45]:

Theorem 6.1 (Stanley). *The labels on the maximal chains in NC_{n+1} are exactly the parking functions of length n , each occurring once.*

The reader wishing to read more about parking functions and noncrossing partitions might want to consult [11], [29], [48], [51], or especially, [45].

7. FREE PROBABILITY.

Our third sighting of the noncrossing partition lattice is in a noncommutative version of probability. Because of the nature of the subject matter, the discussion

in this section is less detailed than in previous ones. Readers wishing to read more about the connection between free probability and noncrossing partitions should probably begin with the excellent survey article by Roland Speicher [42]. We start with a brief discussion of classical probability.

Let X be a random variable having a probability density function $f(x)$. For the reader unfamiliar with probability theory, a good example of a probability density function is a nonnegative continuous function from \mathbf{R} to \mathbf{R} whose integral over the reals is 1. The *expectation* of a function $u(X)$ is then defined as

$$E(u(X)) = \int_{-\infty}^{\infty} u(x)f(x)dx.$$

The first expectations students usually encounter are the mean $\mu = E(X)$ and the variance $\sigma^2 = E((X - \mu)^2)$. In a mathematical probability and statistics course they meet the higher moments $E(X^n)$ as well as the moment generating function

$$M(t) = E(e^{tX}) = \sum_{n \geq 0} E(X^n) \frac{t^n}{n!},$$

which allows the calculation of all of the moments by evaluating a single integral. The coefficients in the moment generating function are called the (classical) *moments* of X . The coefficients of $\log M(t)$ are called the (classical) *cumulants* of X . The main advantage of the cumulants is that they contain the same information as the moments but that the n th cumulant of the sum of two random variables X and Y is the sum of their n th cumulants—provided they are independent.

Before launching into the “noncommutative” version, we should say a brief word about noncommutative geometry in general. Noncommutative geometry is a philosophy whereby standard geometric arguments on topological spaces are converted into algebraic arguments on their commutative C^* -algebras of functions. The motivation comes from mathematical physics and the need to integrate quantum mechanics (which is noncommutative) with classical physics. The main observation is that there is a nice correspondence between “reasonable” topological spaces X and the collections of continuous maps from these spaces to the complex numbers. Each such collection has a structure known as a C^* -algebra. In fact, the correspondence is strong enough that a space X can be recovered from the commutative C^* -algebra to which it gives rise. The philosophy, in short, is to reformulate each concept from classical topology, geometry, calculus of manifolds, and so forth in terms of properties of C^* -algebras and then to turn these equivalent formulations into definitions for a “noncommutative” version of this concept. Even though there are no longer any topological spaces or points or open sets—only C^* -algebras—the classical structures can be used to develop intuition and provide a guide to the types of theorems and results that should be expected from the noncommutative world. To date this program has been remarkably successful. The original book by Alain Connes [18] or the more recent (and shorter) survey articles such as [19] are a good place for an inexperienced reader to begin.

Returning to probability theory, we remark that researchers have defined a noncommutative probability space to be a pair (\mathcal{A}, ϕ) , where \mathcal{A} is a complex-unital-algebra equipped with a unital linear functional ϕ called expectation. There is also a noncommutative version of independence known as “freeness.” Without getting into the details, the combinatorics of noncrossing partitions is very closely involved in the noncommutative version of cumulants. In fact, some researchers

who study free probability have described the passage from the commutative to the noncommutative setting of probability as a transition from the combinatorics of the partition lattice Π_n to the combinatorics of noncrossing partitions NC_n . The article by Roland Speicher [42] is an excellent exposition of this topic. To give just one hint at the underlying argument, we note that new counting problems that involve summing up an old counting problem over all of the possible partitions of that problem into smaller problems of the same type often lead to solutions that contain exponentials (see [47, chap. 5] for a precise development of this theme). Thus the exponential function in the integral defining the moment generating function gets converted into a sum over the elements in the partition lattice. In the noncommutative context, the crossing partitions are prevented from playing a role, so the sum takes place over the noncrossing partition lattice instead.

8. SUMMARY.

As we have seen, the lattice of noncrossing partitions might surface in any situation that involves (1) the symmetric groups, (2) the braid groups, (3) free probability, or (4) the Catalan numbers (say, in conjunction with the combinatorics of trees or the combinatorics of parking functions). They also show up in real hyperplane arrangements, Prüfer codes, quasisymmetric functions, and Hopf algebras, but detailing all of these connections would lead us too far afield. (see [1], [4], [5], [6], [8], [24], [46], [47], or [45] for details.) We conclude with one final illustration: the associahedron.

Associahedron. As we remarked earlier, the Catalan numbers count the number of ways to associate a list of numbers. If we also consider partial associations such as $(123)(45)$, we get a partial ordering of partial associations (where removing a pair of matched parentheses corresponds to moving up in the ordering). This partial ordering has been shown to be the lattice of faces for a convex polytope known as the *associahedron* (or *Stasheff polytope*). For example, the partial associations of four elements form a pentagon (Figure 10). The full associations label the vertices as indicated, and the partial associations label the edges. For example, $((12)34)$ is the partial association corresponds to the edge connecting the vertices labeled $((((12)3)4)$ and $((1(2)3)4)$ since it can be obtained from either one by removing a pair of parentheses. Similarly, the fourteen associations listed in Figure 2 can be identified with the fourteen vertices of the polytope shown in Figure 11.




FIGURE 10. The two dimensional associahedron.

It is the “Morse theory” of this polytope that has a close connection with the noncrossing partition lattice. If we chose a height function (i.e., a linear map from

the Euclidean space containing the polytope onto the reals) so that none of the edges are horizontal (i.e., none of the direction vectors of the edges lie in the kernel of this map), then at each vertex we can count the number of edges pointing “up” and the number pointing “down.” For the pentagon pictured in Figure 10 (with a height function that orthogonally projects onto the arrow shown) there is one vertex with both edges pointing up, three vertices with one up and one down, and one vertex with both edges pointing down. We can tally these results in a vector known as the *h*-vector of the polytope. Thus, the *h*-vector of the pentagon is $(1, 3, 1)$. Although it is not obvious from this definition, this sequence is independent of the height function chosen.




FIGURE 11. The 3-dimensional associahedron.

The rank function for the noncrossing partition lattice NC_3 (i.e., the number of elements it contains at each height) can also be summarized by a vector. As seen in Figure 7, the vector for NC_3 is $(1, 3, 1)$ since there is one element with height zero, three elements with height one, and one element with height two. The recurrence of the vector $(1, 3, 1)$ is not a coincidence. Notice that the associahedron shown in Figure 11 has an *h*-vector $(1, 6, 6, 1)$ that corresponds nicely with the rank function for NC_4 (Figure 5).

Fomin and Zelevinsky have recently defined for each finite crystallographic reflection group a Euclidean polytope known as a *generalized associahedron* [25], [14], and in each case the *h*-vector for the general polytope matches the rank function for the corresponding lattice [4]. These lattices and polytopes, and the observed connections between them, are only a few years old at this point. They are the subject of many ongoing research projects. I am sure that we will be discovering additional remarkable properties of these objects for many years to come.

ACKNOWLEDGMENTS. In 1999 Texas A&M University hired three new assistant professors: myself (geometric group theory), Catherine Huafei Yan (combinatorics), and Ken Dykema (free probability). The discovery that we were all interested in the structure of the noncrossing partition lattice from vastly different perspectives marked the early beginnings of this article. The more proximate cause

was a general-interest talk I gave at the conference celebrating Jim Cannon's sixtieth birthday held in Park City, Utah, in June, 2003. I would like to thank the organizers for their invitation. I would also like to thank Christos Athanasiadis, Tom Brady, Ken Dykema, John Meier, Vic Reiner and Jim Stasheff for comments on early versions of the article. As a final note, a colorful version of this article is available from my homepage (<http://www.math.ucsb.edu/~mccammon/>). The author was partially supported by the National Science Foundation.

REFERENCES

- [1] M. Aigner, Catalan and other numbers: A recurrent theme, in *Algebraic Combinatorics and Computer Science*, Springer Italia, Milan, 2001, pp. 347–390.
- [2] M. Anshelevich, Free stochastic measures via noncrossing partitions, *Adv. Math.* **155** (2000) 154–179.
- [3] ———, Free stochastic measures via noncrossing partitions II, *Pacific J. Math.* **207** (2002) 13–30.
- [4] C. A. Athanasiadis, On a refinement of the generalized Catalan numbers for Weyl groups, *Trans. Amer. Math. Soc.* **357** (2005) 179–196.
- [5] ———, On noncrossing and nonnesting partitions for classical reflection groups, *Electron. J. Combin.* **5** (1998) research paper no. 42, 16 pp. (electronic).
- [6] C. A. Athanasiadis and S. Linusson, A simple bijection for the regions of the Shi arrangement of hyperplanes, *Discrete Math.* **204** (1999) 27–39.
- [7] C. A. Athanasiadis and V. Reiner, Noncrossing partitions for the group D_n , *SIAM Journal on Discrete Mathematics* **18** (2004) 397–417.
- [8] J.-C. Aval and N. Bergeron, Catalan paths and quasi-symmetric functions, *Proc. Amer. Math. Soc.* **131** (2003) 1053–1062.
- [9] D. Bessis, The dual braid monoid, *Ann. Sci. École Norm. Sup.* **36** (2003) 647–683.
- [10] M. Bestvina, Non-positively curved aspects of Artin groups of finite type, *Geom. Topol.* **3** (1999) 269–302.
- [11] P. Biane, Parking functions of types A and B, *Electron. J. Combin.* **9** (2002) 5 pp. (electronic).
- [12] T. Brady, A partial order on the symmetric group and new $K(\pi, 1)$'s for the braid groups, *Adv. Math.* **161** (2001) 20–40.
- [13] T. Brady and C. Watt, $K(\pi, 1)$'s for Artin groups of finite type, *Geom. Dedicata* **94** (2002) 225–250.
- [14] F. Chapoton, S. Fomin, and A. Zelevinsky, Polytopal realizations of generalized associahedra, *Canad. Math. Bull.* **45** (2002) 537–566.
- [15] R. Charney, Geodesic automation and growth functions for Artin groups of finite type, *Math. Ann.* **301** (1995) 307–324.
- [16] R. Charney and M. W. Davis, Finite $K(\pi, 1)$ s for Artin groups, in *Prospects in Topology (Princeton, NJ, 1994)*, Princeton University Press, Princeton, 1995, pp. 110–124.
- [17] A. M. Cohen and D. B. Wales, Linearity of Artin groups of finite type, *Israel J. Math.* **131** (2002) 101–123.
- [18] A. Connes, *Noncommutative Geometry*, Academic Press, San Diego, 1994.
- [19] ———, A short survey of noncommutative geometry, *J. Math. Phys.* **41** (2000) 3832–3866.
- [20] N. Dershowitz and S. Zaks, Ordered trees and noncrossing partitions, *Discrete Math.* **62** (1986) 215–218.
- [21] P. H. Edelman, Chain enumeration and noncrossing partitions, *Discrete Math.* **31** (1980) 171–180.
- [22] ———, Multichains, noncrossing partitions and trees, *Discrete Math.* **40** (1982) 171–179.
- [23] P. H. Edelman and R. Simion, Chains in the lattice of noncrossing partitions, *Discrete Math.* **126** (1994) 107–119.
- [24] R. Ehrenborg, On posets and Hopf algebras, *Adv. Math.* **119** (1996) 1–25.
- [25] S. Fomin and A. Zelevinsky, Y-systems and generalized associahedra, *Ann. of Math.* **158** (2003) 977–1018.
- [26] J. Gilman, W. Menasco, and X.-S. Lin, eds., *Knots, Braids, and Mapping Class Groups—Papers Dedicated to Joan S. Birman*, American Mathematical Society, Providence, 2001.
- [27] I. Grossman and W. Magnus, *Groups and Their Graphs*, Random House, New York, 1964.

- [28] V. Jones and V. S. Sunder, *Introduction to Subfactors*, Cambridge University Press, Cambridge, 1997.
- [29] L. H. Kalikow, Symmetries in trees and parking functions, *Discrete Math.* **256** (2002) 719–741.
- [30] C. Kassel, M. Rosso, and V. Turaev, *Quantum Groups and Knot Invariants*, Société Mathématique de France, Paris, 1997.
- [31] D. Krammer, The braid group B_4 is linear, *Invent. Math.* **142** (2000) 451–486.
- [32] ———, Braid groups are linear, *Ann. of Math.* **155** (2002) 131–156.
- [33] G. Kreweras, Sur les partitions non croisées d'un cycle, *Discrete Math.* **1** (1972) 333–350.
- [34] K. Murasugi and B. I. Kurpita, *A Study of Braids*, Kluwer, Dordrecht, 1999.
- [35] A. Nica and R. Speicher, A “Fourier transform” for multiplicative functions on non-crossing partitions, *J. Algebraic Combin.* **6** (1997) 141–160.
- [36] V. V. Prasolov and A. B. Sossinsky, *Knots, Braids and 3-Manifolds*, American Mathematical Society, Providence, 1997.
- [37] H. Prodinger, A correspondence between ordered trees and noncrossing partitions, *Discrete Math.* **46** (1983) 205–206.
- [38] D. Rolfsen, New developments in the theory of Artin’s braid groups, *Topology Appl.* **127** (2003) 77–90.
- [39] R. Simion, Noncrossing partitions, *Discrete Math.* **217** (2000) 367–409.
- [40] R. Simion and D. Ullman, On the structure of the lattice of noncrossing partitions, *Discrete Math.* **98** (1991) 193–206.
- [41] R. Speicher, Multiplicative functions on the lattice of noncrossing partitions and free convolution, *Math. Ann.* **298** (1994) 611–628.
- [42] ———, Free probability theory and non-crossing partitions, *Sém. Lothar. Combin.* **39** (1997) 38 pp. (electronic).
- [43] ———, Combinatorial theory of the free product with amalgamation and operator-valued free probability theory, *Mem. Amer. Math. Soc.* **132** (1998), x+88.
- [44] R. P. Stanley, Exercises on Catalan numbers and related numbers, available at <http://www-math.mit.edu/~rstan/ec/>.
- [45] ———, Parking functions and noncrossing partitions, *Electron. J. Combin.* **4** (1997) 14 pp. (electronic).
- [46] ———, *Enumerative combinatorics*, vol. 1, Cambridge University Press, Cambridge, 1997.
- [47] ———, *Enumerative combinatorics*, vol. 2, Cambridge University Press, Cambridge, 1999.
- [48] R. P. Stanley and J. Pitman, A polytope related to empirical distributions, plane trees, parking functions, and the associahedron, *Discrete Comput. Geom.* **27** (2002) 603–634.
- [49] D. Voiculescu, Free probability theory: Random matrices and von Neumann algebras, in *Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994)*, Birkhäuser, Basel, 1995, pp. 227–241.
- [50] ———, Lectures on free probability theory, in *Lectures on Probability Theory and Statistics (Saint-Flour, 1998)*, Springer-Verlag, Berlin, 2000, pp. 279–349.
- [51] C. H. Yan, On the enumeration of generalized parking functions, *Congr. Numer.* **147** (2000) 201–209.

JON MCCAMMOND is an associate professor at the University of California, Santa Barbara. Although his research is primarily focused on geometric group theory, he has a more than passing interest in geometric combinatorics (particularly polytopes), theoretical computer science (particularly finite state automata), and differential geometry (particularly symmetric spaces).

DEPT. OF MATH., U. C. SANTA BARBARA, SANTA BARBARA, CA 93106
E-mail address: jon.mccammond@math.ucsb.edu